

We claim:

1. A method comprising:

identifying a source volume containing stored data, wherein the source volume is a virtual volume comprising objects defining a mapping to data in at least one storage device, and wherein the objects are distributed across more than one processor in a virtualization layer between at least one host and the at least one storage device;

generating a snapshot of the source volume that is a point-in-time copy containing state information about a state of the source volume when the snapshot is generated; and

distributing the snapshot of the source volume across the more than one processor in the virtualization layer.

2. The method of claim 1, wherein the method of the snapshot of the source volume does not affect the availability of the source volume to the at least one host.

3. The method of claim 1, wherein the generating further includes:

for each processor that contains objects in the virtual volume, creating an intermediate snapshot of a portion of the virtual volume corresponding to objects contained on the processor.

4. The method of claim 3, wherein the generating further includes:

combining the intermediate snapshots from each processor to generate the snapshot of the source volume.

5. The method of claim 1, further comprising:
creating a change log corresponding to the snapshot; and
storing in the change log changes to the source volume made after the snapshot is generated.

6. The method of claim 5, wherein the change log is a copy on write (COW) change log.

7. The method of claim 5, wherein snapshot cannot be changed after it is generated.

8. A method comprising:

identifying a source volume containing stored data, wherein the source volume is a virtual volume comprising objects defining a mapping to data in at least one storage device, and wherein the objects are distributed across more than one processor in a virtualization layer between at least one host and the at least one storage device;

specifying a change log volume corresponding to the source volume;

generating a snapshot of the source volume that is a point-in-time copy containing state information about a state of the source volume when the snapshot is generated; and

storing, in the change log volume, changes made to the source volume after the snapshot is generated.

9. The method of claim 8, further comprising:

distributing the snapshot of the source volume across the more than one processor in the virtualization layer; and

distributing the change log volume across the more than one processor in the virtualization layer.

10. The method of claim 8, further comprising:

receiving a request for data stored in the source volume;

determining from the change log volume, whether the requested data has changed since the snapshot was generated;

retrieving the requested data from the change log volume, when it is determined that the requested data has changed since the snapshot was generated; and

retrieving the requested data from the source volume, when it is determined that the requested data has not changed since the snapshot was generated.

11. The method of claim 10, further comprising:

retrieving the requested data from the snapshot, when it is determined that the requested data has not changed since the snapshot was generated.

12. The method of claim 8, wherein the generating of the snapshot of the source volume does not affect the availability of the source volume to the at least one host.

13. The method of claim 8, wherein the generating further includes:

for each processor that contains objects in the virtual volume, creating an intermediate snapshot of a portion of the virtual volume corresponding to objects contained on the processor.

14. The method of claim 13, wherein the generating further includes:

combining the intermediate snapshots from each processor to generate the snapshot of the source volume.

15. A system comprising:

a plurality of storage devices storing data corresponding to a host;

a virtualization layer between the host and the plurality of storage devices, the virtualization layer comprising objects defining a mapping to data in the plurality of storage devices; and

a snapshot layer between the host and the virtualization layer, the snapshot layer comprising:

an intermediate snapshot of each object in the virtualization layer, the intermediate snapshot having references to (1) the object in the virtualization layer, (2) an intermediate point-in-time copy of the object, and (3) a change log corresponding to the intermediate point-in-time copy, and

an overall snapshot object containing a reference to each intermediate snapshot.

16. The system of claim 15, wherein each intermediate point-in-time copy contains state information about a state of the corresponding object in the virtualization layer when the snapshot layer is generated.

17. The system of claim 15, wherein each change log stores changes made to the corresponding intermediate point-in-time copy after the snapshot layer is generated.

18. The system of claim 15, wherein the virtualization layer has multiple processors and the objects defining the mapping to data in the plurality of storage devices are distributed across the multiple processors.

19. The system of claim 15, wherein the virtualization layer has multiple processors and the intermediate snapshot objects are distributed across the multiple processors.

20. The system of claim 15, further comprising:
an interface enabling the host to view a point-in-time representation of the data by accessing the overall snapshot object.

21. The system of claim 15, further comprising:
an interface enabling the host to specify when the snapshot layer is created.

22. The system of claim 15, wherein the snapshot layer is created on a periodic basis.

23. A system comprising:

means for identifying a source volume containing stored data, wherein the source volume is a virtual volume comprising objects defining a mapping to data in at least one storage device, and wherein the objects are distributed across more than one processor in a virtualization layer between at least one host and the at least one storage device;

means for generating a snapshot of the source volume that is a point-in-time copy containing state information about a state of the source volume when the snapshot is generated; and

means for distributing the snapshot of the source volume across the more than one processor in the virtualization layer.

24. The system of claim 23, wherein the generating means further includes:

means for creating an intermediate snapshot of a portion of the virtual volume corresponding to objects contained on each processor that contains objects in the virtual volume.

25. The system of claim 24, wherein the generating means further includes:

means for combining the intermediate snapshots from each processor to generate the snapshot of the source volume.

26. The system of claim 25, further comprising:

means for creating a change log corresponding to the snapshot; and

means for storing in the change log changes to the source volume made after the snapshot is generated.

27. A computer-readable medium containing code for directing a processor to perform a method for creating a copy of stored data, the method comprising:

identifying a source volume containing stored data, wherein the source volume is a virtual volume comprising objects defining a mapping to data in at least one storage device, and wherein the objects are distributed across more than one processor in a virtualization layer between at least one host and the at least one storage device;

generating a snapshot of the source volume that is a point-in-time copy containing state information about a state of the source volume when the snapshot is generated; and

distributing the snapshot of the source volume across the more than one processor in the virtualization layer.

28. The computer-readable medium of claim 27, wherein the method further includes:

for each processor that contains objects in the virtual volume, creating an intermediate snapshot of a portion of the virtual volume corresponding to objects contained on the processor.

29. The computer-readable medium of claim 27, wherein the method further includes:

combining the intermediate snapshots from each processor to generate the snapshot of the source volume.

30. The computer-readable medium of claim 27, wherein the method further includes:
- creating a change log corresponding to the snapshot; and
 - storing in the change log changes to the source volume made after the snapshot is generated.